

REMARKS

This communication is responsive to the Office Action dated June 2, 2006. A petition for a three month extension of the time period for responding to the Action is separately filed herewith. In the amendment, claims 55, 56, 63, 64, 71, 72, 79 and 80 have been canceled without prejudice, and claims 54, 62, 70, 78, 86, 88, 90 and 91 have been amended, such that claims 54, 57-62, 65-70, 73-78 and 81-92 remain pending in this application. These amendments add no new matter. Applicant respectfully requests reconsideration of the pending claims in view of these amendments and the following remarks.

Applicant notes with appreciation the courtesy of Examiner Shingles in extending the opportunity for a telephone interview with the undersigned representative regarding the outstanding Office Action, on November 29, 2006.

In the interview, the specification and amended claims were discussed, as well as the rejection of the claims under 35 U.S.C. § 102 as being anticipated. Applicant described features relating to the storage domains that were understood as being distinct from anything disclosed by the prior art, including the reference relied upon in the Action. Applicant submits that the claims presently in the application reflect these distinctions.

Claims 54-92 were rejected under 35 U.S.C. § 102(e) as being anticipated by U.S. Patent No. 6,597,956 to Aziz et al. (“Aziz”). This rejection is traversed.

Claim 54 has been amended and now recites: *[a] method for managing a storage area network (SAN), the method comprising:*

defining storage domains respectively having associated configurable storage management properties that are separate from individual physical capabilities of physical storage resources available through the SAN, wherein a first set of storage management properties is associated with a first storage domain and a second set of storage management properties is associated with a second storage

domain, with the first set of storage management properties being different from the second set of storage management properties; accommodating the creation of logical volumes configurable for presentation to hosts through the SAN; and allocating the logical volumes to hosts in the context of the storage domains, wherein allocating a first logical volume to a first host in the context of the first storage domain entails the provision of storage resources according to the first set of storage management properties and allocating a second volume to a second host in the context of the second storage domain entails the provision of storage resources according to the second set of storage management properties wherein presentation of logical volumes to hosts accommodates storage resource access by hosts without requiring hosts to be configured according to the requirements of the physical storage resources, and wherein said allocation of the first logical volume to the first host entails the provision of storage resources according to the first set of storage management properties and said allocation of the second logical volume to the second host subject entails the provision of storage resources according to the second set of storage management properties even though the first logical volume and second logical volume are created from physical storage resources in common.

Applicant's claimed invention accommodates improved management of storage resources in the SAN environment by providing distinct storage domains that are separate from the individual physical capabilities of physical storage resources available through the SAN, and which respectively have associated configurable storage management properties. Moreover, with the claimed invention, logical volumes in different storage domains (with respective storage management properties) may be created from physical storage resources in common. This is because the storage management properties are dictated by the context of allocation (*i.e.*, which storage domain), rather than according to the traditional division and assignment of physical storage having particular properties. These features allow the volumes to have different storage

management properties according to the storage domain in which their allocation to a host is made. This allows different storage management properties even though the volumes are created from physical storage resources in common.

Zoning (*e.g.*, as described in Aziz) divides storage so that particular pieces can be managed and assigned as desired (*e.g.*, dividing storage into zone#1, zone#2, etc.). With these zones, there is a mapping (*e.g.*, by port) of hosts to their particular designated zone. This division provides “pieces” of storage that have respective features and characteristics.

By contrast, with Applicant’s claimed invention, the storage domains are separate from the physical storage resources. Physical storage resources in common (*e.g.*, the same “piece” of storage) can be used to create logical volumes that have completely different storage management properties. These different properties are imposed not by ensuring that respective volumes correspond to a particular zone or division, but rather by virtue of their allocation in the context of different storage domains.

Thus, even if the same volume is separately allocated to two hosts, the context of allocation – the storage domain – dictates the properties of how such storage is respectively provided for the two hosts. And such is the case even though the logical volumes are created from physical storage resources in common.

Applicant submits that claim 54 is neither disclosed nor suggested by Aziz. Aziz discloses a method and apparatus for controlling an extensible computing system. A “Virtual Server Farm” (VSF) is created from a computing grid which is physically constructed and logically divided into the VSFs for access by various organizations on demand. (Aziz, Abstract). While Aziz may have some advantages in terms of managing these “VSFs”, there is clearly no disclosure of the features of Applicant’s claimed invention. Rather, Aziz merely appears to

invoke typical zoning for storage management, wherein storage resources are divided and then access to the divisions is controlled in some fashion. For example, as stated in Aziz:

“[a]t any given point in time, a computing element in the computing grid, such as CPU1 of FIG. 2, is only connected to the set of VLANs and the SAN zone(s) associated with a single VSF. A VSF typically is not shared among different organizations. The subset of storage on the SAN that belongs to a single SAN zone, and the set of VLANs associated with it and the computing elements on these VLANs define a VSF.

By controlling the membership of a VLAN and the membership of a SAN zone, control plane enforces a logical partitioning of the computing grid into multiple VSFs. Members of one VSF cannot access the computing or storage resources of another VSF. Such access restrictions are enforced at the hardware level by the VLAN switches, and by port-level access control mechanisms (e.g., zoning) of SAN hardware such as Fibre Channel switches and edge devices such as SCSI to Fibre Channel bridging hardware. Computing elements that form part of the computing grid are not physically connected to the control ports or interfaces of the VLAN switches and the SAN switches, and therefore cannot control the membership of the VLANs or SAN zones. Accordingly, the computing elements of the computing grid cannot access computing elements not located in the VSF in which they are contained.”

(Aziz, at 7:18-42). In this passage (and throughout), Aziz describes partitioning of resources according to zones. Access is then controlled such as by “port-level access control mechanisms.” There is no description of separate storage domains that define storage management properties, as claimed by Applicant.

Additionally, any “allocation” that is described in Aziz is clearly distinct from the allocation claimed by Applicant. Aziz is devoid of any particulars with regard to the allocation of logical volumes, even in a general sense. Aziz uses the term “allocate”, but this usage appears to be for allocating processing resources, such as a CPU. For example, in column 10, lines 40-50, relied upon by the Examiner as disclosing certain features of the claimed allocation, Aziz merely states that:

“In response, the control plane assigns or allocates CPU A as the load balancer/firewall, and allocates CPUs B and C as the Web servers. CPU A is logically placed in SAN Zone 1, and pointed to a bootable partition on a disk that

contains dedicated load balancing/firewalling software. The term "pointed to" is used for convenience and is intended to indicate that CPU A is given, by any means, information sufficient to enable CPU A to obtain or locate appropriate software that it needs to operate. Placement of CPU A in SAN Zone 1 enables CPU A to obtain resources from disks that are controlled by the SAN of that SAN Zone."

(Aziz, at 10:40-50). First, this passage clearly does not describe allocating logical volumes to hosts in the context of the storage domains as claimed by Applicant. In any event, to the extent that some subsequent allocation may be presumed, the allocation could only be within the assigned zone. To repeat what Aziz states, "[p]lacement of CPU A in SAN Zone 1 enables CPU A to obtain resources from disks that are controlled by the SAN of that SAN Zone." (Aziz, at 10:48-50).

The zoning in Aziz clearly does not disclose or suggest the features "*wherein said allocation of the first logical volume to the first host entails the provision of storage resources according to the first set of storage management properties and said allocation of the second logical volume to the second host subject entails the provision of storage resources according to the second set of storage management properties even though the first logical volume and second logical volume are created from physical storage resources in common,*" as claimed by Applicant.

SAN Zone 1 of Aziz has certain assigned physical storage resources. If there is another SAN Zone (e.g., 2, or 3, etc.), it would have its own separate physical storage resources. Accordingly, in Aziz there is quite clearly no allocation of first and second logical volumes in the context of first and second storage domains to provide storage according to different storage management properties, even though the first and second logical volumes are created from physical storage resources in common, as claimed by Applicant. Even if one allocates multiple volumes within a given Zone, the storage management properties of that allocation would be

dictated by that given Zone. On the other hand, if multiple Zones are involved, they could conceivably have different properties, but then "*physical storage resources in common*" would not be used in the respective allocations. In either case, the features recited in independent claim 54 are clearly absent from Aziz.

For reasons similar to those provided regarding claim 54, independent claims 62, 70 and 78 are also neither disclosed or suggested by Aziz, and Applicant requests reconsideration and withdrawal of the rejection of those claims. Claims 57-61, 65-69, 73-77 and 81-92 depend from these independent claims and thus are distinct from Aziz for their incorporation of the features recited in the independent claims, as well as for the patentably distinct features separately recited therein, and previously made of record.

Accordingly, Applicant requests reconsideration and withdrawal of the rejection of the claims under 35 U.S.C. § 102(e) as being anticipated by Aziz.

For the foregoing reasons, reconsideration and allowance of the claims which remain in this application are solicited. If any further issues remain, the Examiner is invited to telephone the undersigned to resolve them.

Dated:

December 4, 2006

Respectfully submitted,

By _____
Christopher M. Tobin

Registration No. 40,290
RADER, FISHMAN & GRAUER PLLC
1233 20th Street, N.W.
Suite 501
Washington, DC 20036
(202) 955-8779
Attorney for Applicant